To appear in *Statistics: A Journal of Theoretical and Applied Statistics*
Vol. 00, No. 00, Month 201X, 1–18

arXiv:1401.7001v1 [math.ST] 27 Jan 2014

A modified χ^2-test for uplift models with applications in marketing performance measurement

René Michel*, Igor Schnakenburg and Tobias von Martens

Altran GmbH & Co. KG, Schillerstraße 20, 60313 Frankfurt, Germany

(*v1.0 sent January 2014*)

Uplift, essentially being the difference between two probabilities, is a central number in marketing performance measurement. A frequent question in applications is whether the uplifts of two campaigns are significantly different. In this article we present a new χ^2-statistic which allows to answer this question by performing a statistical test. We show that this statistic is asymptotically χ^2-distributed and demonstrate its application in a real life example. By running simulations with this new and alternative approaches, we find our suggested test to exhibit a better decisive power.

Keywords: net scoring; uplift modelling; χ^2-test; campaign management; performance measurement

AMS Subject Classification: 62F03

1. Introduction

Nowadays companies do not wait for customers to contact them and show their interest in a certain product. They rather try to initiate the contact themselves. This can be done via channels such as mail, email or phone call. The company-initiated, personalised targeting of a fixed group of persons with a specific topic is usually referred to as a campaign. These campaigns are part of the so called below-the-line marketing since only the targeted group is contacted. An integral part of the campaign management process is to measure, whether a campaign was successful in order to be able to optimize the company's campaign portfolio. Since customer behaviour is subject to random fluctuations, statistical methods are needed to distinguish between "true" effects and success by chance.

One can measure campaign effectiveness only on the basis of differences between response rates in a target group and a structurally similar group that has not been addressed with the campaign (control group). This difference is usually referred to as uplift, see [1]. Only in case it is positive, a campaign has generated added value and may therefore be regarded as effective, see [2]. An overview of the statistics in campaign performance measurement is given in Chapter 8 of [3].

In this paper we present a statistical method of comparing success of two campaigns, or more precisely, of comparing their uplifts. We do this by introducing χ^2_{net}, a modified version of the classical χ^2-statistic. With this new statistic and by showing that it asymp-

*Corresponding author. Email: rene.michel@altran.com

totically follows a χ^2-distribution, we are able to discover significant differences between the uplifts of two campaigns. We will also show its application to real data examples.

Up to now, only few studies cover this area of research. They primarily deal with the context of net scoring in which our statistic can also be used. Net scoring is a generalisation of classical scoring procedures in that it does not predict the probability of a desired customer behaviour, but instead predicts the increase of this probability through targeting. Some algorithms used in this context apply statistics in order to compare the uplifts of two campaigns. The most prominent example is [4]. Here, an F-statistic is introduced which tackles this problem. Another example is [5] where an ancestor of our statistic χ^2_{net} is introduced. Both papers primarily show the statistics and apply them as parts of complex algorithms but do not investigate their statistical background. In this paper, we focus on the mathematical background. We construct a sound statistical testing procedure for the comparison of campaign success which includes the computation of asymptotical distributions.

The alternatives to our approach are based on variance analysis. The first is the above mentioned F-statistic, the second would be the use of contrasts. We will also introduce these methods and compare them to our new approach, mainly by simulations.

We begin by summarizing the basics of campaign performance measurement, especially control groups, in Section 2. After that, we present the test scenarios which we investigate in this paper. Section 4, being the main section of our article, introduces the χ^2_{net}-statistic and proves its asymptotic χ^2-distribution. The alternative approaches are summarized in Section 5. They are compared in Section 6 by means of a simulation study. In Section 7, we show an application of the statistics to a real data example. A discussion and an outlook to yet unsolved problems in Section 8 conclude this paper.

2. Control groups

In order to isolate the impact of campaigns, control groups are necessary. A control group is a structurally identical group of customers that is not targeted and, hence, reflects customer behaviour without campaign impact. In other words, the control group has to be representative of the target group. This is usually ensured by a random selection. Differences in customer behaviour between target and control group may be represented by differences in the respective response rates, see Chapter 8 of [3].

By campaigning, a company usually wishes to induce a certain response behaviour on the targeted people. If a person shows this behaviour within a certain time period after the contact, this is regarded as a response. The number of responses in relation to the number of targeted persons is called response rate and is our means of comparing target and control group

$$\text{response rate} = \frac{\text{number of responses}}{\text{number of targeted persons}}$$

Classical examples of responses are product purchases, the acquisition of a new customer, appointments with sales personnel or the cancellation of business relationships (churn events, in which case a low rate is considered "good").

The response rate in the group of people that have been targeted does not suffice to estimate the success of the campaign under consideration. Campaign effectiveness is measured by comparing it to the rate within the control group. This difference is usually referred to as uplift.

$$\text{uplift} = \text{response rate within target group} - \text{response rate within control group}$$

Only in the case it is positive, a campaign has generated added value and may therefore be regarded effective. The response rate within the control group is often referred to as random noise.

The uplift multiplied by the size of the target group, describes the additional responses gained by the campaign and thus the net impact of that campaign. Examples are additional product sales, additional customers won, additional appointments made or prevented churn events.

This method of measuring the campaign effect is subject to random fluctuations. One classical problem is the question of whether the uplift is different from 0 only by random fluctuations or if a "true" impact is present. This problem can be solved by Fisher's exact test or the classical χ^2-test for homogenity for the comparison of response rates within target and control group, see Sections 4.1 and 4.2 of [6] and Section 8.3 of [3]. Typically, the null hypothesis assumes no impact, i.e. uplift $= 0$.

The problem we wish to discuss in this paper, however, is more complicated. We assume that we have two campaigns and we want to compare their success to each other, i.e. we want to know if their uplifts are significantly different from each other.

3. Test scenarios

In describing our test scenarios another rate will play an important role. Since the validity of our measuring method depends on the target and control groups sizes we define the target-control rate as follows:

$$\text{target-control rate} = \frac{\text{number of persons in target group}}{\text{number of persons in control group}}$$

In Section 4, we will present our test statistic. But first we will explain the two scenarios to which this statistic can be applied.

- Scenario 1:
  Imagine two different campaigns which should be compared regarding their uplift. In order to provide useful results in the marketing context, they should be comparable in some way, e.g. have been carried out at roughly the same time or have advertised a similar product.
- Scenario 2:
  Imagine a campaign with target and control group. Within this campaign there are two different groups of interest, for example men and women or people below and above the age of 40. We are interested in whether the uplifts within both subgroups differ significantly from each other.

Mathematically, one could argue that Scenario 2 is part of Scenario 1. This is true. Assuming a representative control group of the main campaign, the splitting criterion should separate target and control group in an even way. Especially the target-control rate in both subgroups should roughly be equal. In Scenario 1, we make no assumption on the target-control rate. Thus, the case of equal target-control rate in both campaigns is a special case and hence included in Scenario 1. However, we distinguish between these two scenarios, since they represent two different cases from the perspective of campaign

management. In the first case, "independent" campaigns are compared, in the latter case only one campaign is subject to the investigation and is artificially split. Scenario 2 is mainly the way net scoring is done when one uses scoring procedures like decision trees.

4. The χ^2_{net}-statistic

In order to introduce our new χ^2_{net}-statistic, we firstly formalise the setup. We start by assuming two campaigns, each containing target and control group either resulting from one campaign, split into two subgroups by some criterion, or two separate campaigns). In order to ease the notation, we refer to those individual campaigns from now on as subgroups (still covering both Scenarios 1 and 2). If we unify target and control groups, we assume that we have n observations in the unified target group and k observations in the control unified group. We further assume that subgroup 1 contains n_1 target observations and k_1 control observations, subgroup 2 n_2 and k_2 observations respectively. By $a_{G,SG}$, we denote the number of responses in the respective group where $G = T, C$ for target and control group and $SG = 1, 2$ or missing for the first, the second subgroup or the overall group. In tabular form, this reads as

target	response	no response	total
subgroup 1	$a_{T,1}$	$n_1 - a_{T,1}$	n_1
subgroup 2	$a_{T,2}$	$n_2 - a_{T,2}$	n_2
total	a_T	$n - a_T$	n

(1)

for the target group and for the control group

control	response	no response	total
subgroup 1	$a_{C,1}$	$k_1 - a_{C,1}$	k_1
subgroup 2	$a_{C,2}$	$k_2 - a_{C,2}$	k_2
total	a_C	$k - a_C$	k

(2)

We now define the additional responses (or uplift-based responses) of the subgroups of our campaign as

$$l_i = a_{T,i} - \frac{n_i}{k_i} a_{C,i}, \qquad i = 1, 2 \tag{3}$$

by simply scaling the responses of the control group to the size of the target group. Note that this definition cannot be used without index, i.e. it cannot be used to define the uplift-based response of the overall campaign, since we do not assume equal target-control rates. Unequal target control rates result in a unified control group being structurally different from the unified target group and, thus, the overall control group not representative of the overall target group. For now, we define the uplift-based responses as the sum of the individual uplift-based responses, i.e., $l := l_1 + l_2$. In tabular form this reads

uplift	additional responses	no additional responses	total
subgroup 1	l_1	$n_1 - l_1$	n_1
subgroup 2	l_2	$n_2 - l_2$	n_2
total	l	$n - l$	n

We next make the following model assumption: $a_{G,SG}$, $G = T, C$, $SG = 1, 2$ follow a binomial distribution with a response probability $0 < p_{G,SG} < 1$. We also assume that the responses in the four target and control subgroups are independent of each other.

In the following, we want to investigate the hypothesis

$$p_{T,1} - p_{C,1} = p_{T,2} - p_{C,2} \tag{4}$$

i.e. that the uplift is the same in both subgroups.

For that, we compute the expectation and the variance of the uplift-based responses by the known expectation and variance formula for a binomial distributed random variable, see Section 3.2 of [7] and the standard calculation rules for expectation and variance, see Sections 2.3 and 2.4 of [7]

$$\mathrm{E}(l_i) = n_i p_{T,i} - n_i p_{C,i} = n_i(p_{T,i} - p_{C,i}), \qquad i = 1, 2 \tag{5}$$

and

$$\mathrm{Var}(l_i) = n_i p_{T,i}(1 - p_{T,i}) + \frac{n_i^2}{k_i} p_{C,i}(1 - p_{C,i}), \qquad i = 1, 2 \tag{6}$$

We estimate $p_{T,i}$ by $\hat{p}_{T,i} := \frac{a_{T,i}}{n_i}$ and $p_{C,i}$ by $\hat{p}_{C,i} := \frac{a_{C,i}}{k_i}$, $i = 1, 2$ which are unbiased estimators and converge to the probabilties by the law of large numbers, see Section 4.1 of [7].

Next, we estimate the response propabilities of the unified target and control groups. By elementary stochastic considerations we find the probability p_T of the unified target group to be

$$p_T = \frac{n_1}{n} p_{T,1} + \frac{n_2}{n} p_{T,2} \tag{7}$$

and it can be estimated by

$$\hat{p}_T := \frac{a_T}{n} = \frac{n_1}{n} \hat{p}_{T,1} + \frac{n_2}{n} \hat{p}_{T,2} \tag{8}$$

For the control group, we could make an analogous defintion which however leads to biased results in Scenario 1. An analogous definition for the control group would however lead to biased results considering the setup of Scenario 1 where we cannot make the assumption of equal target-control group rates. We therefore define

$$p_C := \frac{n_1}{n} p_{C,1} + \frac{n_2}{n} p_{C,2}, \tag{9}$$

where the "natural" weighting factors k, k_1 and k_2 have been replaced by n, n_1 and n_2. An unbiased estimator is

$$\hat{p}_C := \frac{n_1}{n} \hat{p}_{C,1} + \frac{n_2}{n} \hat{p}_{C,2} \tag{10}$$

This definition of p_C by the target group weighted control group response probabilities is crucial, as we will explain later. A discussion of the case $\hat{p}_C := \frac{a_C}{k}$ will be given in

Remark 4.2.
We put

$$\hat{e}_i = n_i \left(\hat{p}_T - \hat{p}_C\right) \tag{11}$$

as an estimator of the expectation of the additional responses in each individual campaign which is only valid under the null hypothesis (4) (note the missing indices on the right hand side). Additionally, we define

$$\hat{v}_i = n_i \hat{p}_{T,i}(1-\hat{p}_{T,i}) + \frac{n_i^2}{k_i}\hat{p}_{C,i}(1-\hat{p}_{C,i}) \tag{12}$$

as an unbiased estimator of the variances of l_i.

In order to test the null hypothesis (4), we define the net χ^2-statistic

$$\chi^2_{\text{net}} := \left(\frac{(l_1-\hat{e}_1)^2}{\hat{v}_1} + \frac{(l_2-\hat{e}_2)^2}{\hat{v}_2}\right) \cdot \frac{1}{\hat{w}_n \hat{f}_n} \tag{13}$$

with the norming terms

$$\begin{aligned}\hat{w}_n = {} & \frac{n_2}{n}\left[\hat{p}_{T,1}(1-\hat{p}_{T,1}) + \frac{n_1}{k_1}\hat{p}_{C,1}(1-\hat{p}_{C,1})\right] \\ & + \frac{n_1}{n}\left[\hat{p}_{T,2}(1-\hat{p}_{T,2}) + \frac{n_2}{k_2}\hat{p}_{C,2}(1-\hat{p}_{C,2})\right]\end{aligned} \tag{14}$$

$$\hat{f}_n = \frac{\frac{n_2}{n}}{\hat{p}_{T,1}(1-\hat{p}_{T,1}) + \frac{n_1}{k_1}\hat{p}_{C,1}(1-\hat{p}_{C,1})} + \frac{\frac{n_1}{n}}{\hat{p}_{T,2}(1-\hat{p}_{T,2}) + \frac{n_2}{k_2}\hat{p}_{C,2}(1-\hat{p}_{C,2})} \tag{15}$$

Since $\hat{e}_i$ estimate the expectation of l_i within the unified subgroups (and hence their sum the additional responses for the the overall group), it is easy heuristically to see with the law of large numbers that χ^2_{net} will be close to 0 if both subgroups have the same uplift:

$$\begin{aligned} l_1 - \hat{e}_1 &= a_{T,1} - \frac{n_1}{k_1}a_{C,1} - n_1\left(\hat{p}_T - \hat{p}_C\right) \\ &\approx n_1 p_{T,1} - n_1 p_{C,1} - n_1\left(\frac{n_1}{n}p_{T,1} + \frac{n_2}{n}p_{T,2} - \frac{n_1}{n}p_{C,1} - \frac{n_2}{n}p_{C,2}\right) \\ &= \left(n_1 - \frac{n_1^2}{n}\right)p_{T,1} - \left(n_1 - \frac{n_1^2}{n}\right)p_{C,1} - \frac{n_1 n_2}{n}p_{T,2} + \frac{n_1 n_2}{n}p_{C,2} \\ &= \frac{n_1 n_2}{n}\left[(p_{T,1}-p_{C,1}) - (p_{T,2}-p_{C,2})\right] \overset{(4)}{=} 0 \end{aligned}$$

Analogously, $l_2 - \hat{e}_2 \approx 0$ under the null hypothesis. The terms $\hat{v}_i$, $\hat{w}_n$ and $\hat{f}_n$ ensure the χ^2-distribution of χ^2_{net} when the null hypothesis holds.

Note from the last line that χ^2_{net} scales with a factor of $\frac{n_1 n_2}{n}$ when the uplifts in the subgroups are different. In order to be able to compute p-values, we now want to prove that under the null hypothesis (4), χ^2_{net} follows asymptotically a χ^2-distribution with one

degree of freedom. The proof follows the principles of [8], pp. 446, for the special case of a 2×2 contingency table.

In order to show this convergence, we have to introduce some regularity conditions. Let n_1, n_2, k_1, k_2 and k depend on n. Also suppose that

$$\lim_{n\to\infty} \frac{n_1}{n} = s > 0, \quad \lim_{n\to\infty} \frac{n_1}{k_1} = t_1 > 0, \quad \lim_{n\to\infty} \frac{n_2}{k_2} = t_2 > 0 \tag{16}$$

which ensure that the group sizes increase in a "regular" manner, which is a common assumption in such cases. If $t_1 = t_2$, representativity of the unified control group to the target group in the limit is ensured. This is the setup for using χ^2_{net} in Scenario 2 of Section 3 (splitting one large group into subgroups).

Assumptions (16) imply the following convergence

$$\lim_{n\to\infty} \frac{n_2}{n} = \lim_{n\to\infty} \frac{n - n_1}{n} = 1 - \lim_{n\to\infty} \frac{n_1}{n} = 1 - s \tag{17}$$

Regarding the variances, we find the following implications

$$\begin{aligned} \lim_{n\to\infty} \frac{\text{Var}(l_i)}{\hat{v}_i} &= \lim_{n\to\infty} \frac{n_i p_{T,i}(1-p_{T,i}) + \frac{n_i^2}{k_i} p_{C,i}(1-p_{C,i})}{n_i \hat{p}_{T,i}(1-\hat{p}_{T,i}) + \frac{n_i^2}{k_i} \hat{p}_{C,i}(1-\hat{p}_{C,i})} \\ &= \lim_{n\to\infty} \frac{p_{T,i}(1-p_{T,i}) + \frac{n_i}{k_i} p_{C,i}(1-p_{C,i})}{\hat{p}_{T,i}(1-\hat{p}_{T,i}) + \frac{n_i}{k_i} \hat{p}_{C,i}(1-\hat{p}_{C,i})} = 1 \end{aligned} \tag{18}$$

by the above convergence of the estimators

$$\lim_{n\to\infty} \hat{p}_{G,SG} = p_{G,SG}, \qquad G = T, C, \qquad SG = 1, 2 \tag{19}$$

and regularity conditions (16).

Remark that with analogous arguments as in (18) one gets

$$\lim_{n\to\infty} \frac{\frac{1}{\hat{v}_1} + \frac{1}{\hat{v}_2}}{\frac{1}{\text{Var}(l_1)} + \frac{1}{\text{Var}(l_2)}} = 1 \tag{20}$$

which we need later. Finally, we get

$$\begin{aligned} \lim_{n\to\infty} \hat{w}_n &= (1-s)\left[p_{T,1}(1-p_{T,1}) + t_1 p_{C,1}(1-p_{C,1})\right] \\ &\quad + s\left[p_{T,2}(1-p_{T,2}) + t_2 p_{C,2}(1-p_{C,2})\right] =: w \end{aligned} \tag{21}$$

$$\begin{aligned} \lim_{n\to\infty} \hat{f}_n &= \frac{1-s}{p_{T,1}(1-p_{T,1}) + t_1 p_{C,1}(1-p_{C,1})} \\ &\quad + \frac{s}{p_{T,2}(1-p_{T,2}) + t_2 p_{C,2}(1-p_{C,2})} =: f \end{aligned} \tag{22}$$

We can now state our central theorem and show its proof.

Theorem 4.1. *Under the null hypothesis (4), the regularity conditions (16) and* $p_{T,1}, p_{T,2}, p_{C,1}, p_{C,2} \notin \{0, 1\}$, χ^2_{net} *follows for* $n \to \infty$ *a* χ^2*-distribution with one degree of freedom.*

Proof. Using the definitions of l_1, l_2, $\hat{e}_1$ and $\hat{e}_2$ one can see by elementary calculations that

$$l_2 - \hat{e}_2 = \frac{n_1}{n} a_{T,2} - \frac{n_1 n_2}{n k_2} a_{C,2} - \frac{n_2}{n} a_{T,1} + \frac{n_1 n_2}{n k_1} a_{C,1} = -\left(l_1 - \hat{e}_1\right) \tag{23}$$

holds. Thus, putting this result into (13) we have (remark the index 1 in the second term)

$$\hat{f}_n \hat{w}_n \chi^2_{\text{net}} = \frac{(l_1 - \hat{e}_1)^2}{\hat{v}_1} + \frac{(l_1 - \hat{e}_1)^2}{\hat{v}_2} \tag{24}$$

In a first step we will show

$$\sqrt{\frac{n}{n_1 n_2}} (l_1 - \hat{e}_1) \tag{25}$$

is asymptotically normally distributed with a mean of 0 and variance w from (21). In the second step we will put the pieces together and prove that χ^2_{net} follows a χ^2-distribution with one degree of freedom.

In order to see that (25) is normally distributed, we add zeros to (23) and find the form

$$\begin{aligned} l_1 - \hat{e}_1 = {} & \frac{n_2}{n} \left(a_{T,1} - n_1 p_{T,1}\right) - \frac{n_1 n_2}{n k_1} \left(a_{C,1} - k_1 p_{C,1}\right) - \frac{n_1}{n} \left(a_{T,2} - n_2 p_{T,2}\right) \\ & + \frac{n_1 n_2}{n k_2} \left(a_{C,2} - k_2 p_{C,2}\right) + \frac{n_1 n_2}{n} \underbrace{\left[\left(p_{T,1} - p_{C,1}\right) - \left(p_{T,2} - p_{C,2}\right)\right]}_{\overset{(4)}{=} 0} \end{aligned} \tag{26}$$

which simplifies due to the null hypothesis (4) of equal uplifts in both subgroups.

Thus,

$$\begin{aligned} \sqrt{\frac{n}{n_1 n_2}} \left(l_1 - \hat{e}_1\right) = {} & \underbrace{\frac{n_2}{n} \sqrt{\frac{n}{n_1 n_2}} \sqrt{n_1 p_{T,1} \left(1 - p_{T,1}\right)} \underbrace{\frac{a_{T,1} - n_1 p_{T,1}}{\sqrt{n_1 p_{T,1} \left(1 - p_{T,1}\right)}}}_{=: F_{T,1}}}_{=: H_{T,1}} \\ & \underbrace{- \frac{n_1 n_2}{n k_1} \sqrt{\frac{n}{n_1 n_2}} \sqrt{k_1 p_{C,1} \left(1 - p_{C,1}\right)} \underbrace{\frac{a_{C,1} - k_1 p_{C,1}}{\sqrt{k_1 p_{C,1} \left(1 - p_{C,1}\right)}}}_{=: F_{C,1}}}_{=: H_{C,1}} \\ & \underbrace{- \frac{n_1}{n} \sqrt{\frac{n}{n_1 n_2}} \sqrt{n_2 p_{T,2} \left(1 - p_{T,2}\right)} \underbrace{\frac{a_{T,2} - n_2 p_{T,2}}{\sqrt{n_2 p_{T,2} \left(1 - p_{T,2}\right)}}}_{=: F_{T,2}}}_{=: H_{T,2}} \\ & \underbrace{+ \frac{n_1 n_2}{n k_2} \sqrt{\frac{n}{n_1 n_2}} \sqrt{k_2 p_{C,2} \left(1 - p_{C,2}\right)} \underbrace{\frac{a_{C,2} - k_2 p_{C,2}}{\sqrt{k_2 p_{C,2} \left(1 - p_{C,2}\right)}}}_{=: F_{C,2}}}_{=: H_{C,2}} \end{aligned}$$

Since $a_{G,SG}$ follow a binomial distribution, the $F_{G,SG}$ converge to a standard normal distribution by the Central Limit theorem, see Section 4.2 of [7]. With the convergencies

$$\begin{aligned}
\frac{n_2}{n}\sqrt{\frac{n}{n_1 n_2}}\sqrt{n_1} &= \sqrt{\frac{n_2}{n}} \to_{n\to\infty} \sqrt{1-s} \\
\frac{n_1 n_2}{n k_1}\sqrt{\frac{n}{n_1 n_2}}\sqrt{k_1} &= \sqrt{\frac{n_1}{k_1}}\cdot\sqrt{\frac{n_2}{n}} \to_{n\to\infty} \sqrt{t_1(1-s)} \\
\frac{n_1}{n}\sqrt{\frac{n}{n_1 n_2}}\sqrt{n_2} &= \sqrt{\frac{n_1}{n}} \to_{n\to\infty} \sqrt{s} \\
\frac{n_1 n_2}{n k_2}\sqrt{\frac{n}{n_1 n_2}}\sqrt{k_2} &= \sqrt{\frac{n_2}{k_2}}\sqrt{\frac{n_1}{n}} \to_{n\to\infty} = \sqrt{t_2 s}
\end{aligned}$$

which follow from regularity conditions (16) and (17) we see that

$$\begin{aligned}
H_{T,1} &\to_{n\to\infty} N\left(0,(1-s)p_{T,1}\left(1-p_{T,1}\right)\right) \\
H_{C,1} &\to_{n\to\infty} N\left(0,t_1(1-s)p_{C,1}\left(1-p_{C,1}\right)\right) \\
H_{T,2} &\to_{n\to\infty} N\left(0,sp_{T,2}\left(1-p_{T,2}\right)\right) \\
H_{C,2} &\to_{n\to\infty} N\left(0,t_2 sp_{C,2}\left(1-p_{C,2}\right)\right)
\end{aligned}$$

where $N(\mu,\sigma^2)$ denotes the normal distribution with mean μ and variance σ^2. Since

$$(1-s)\left[p_{T,1}\left(1-p_{T,1}\right)+t_1 p_{C,1}\left(1-p_{C,1}\right)\right]+s\left[p_{T,2}\left(1-p_{T,2}\right)+t_2 p_{C,2}\left(1-p_{C,2}\right)\right]=w,$$

the independence of the $H_{G,SG}$ (which follows from the independence of $a_{G,SG}$) and the convolution theorem for the normal distribution (i.e. the sum of two independent normal random variables is also normal), we conclude that $\sqrt{\frac{n}{n_1 n_2}}\left(l_1-\hat{e}_1\right)$ converges under the null hypothesis (4) to a normal distribution with a mean of 0 and variance w, which finalizes the first step of our proof.

We begin the second step by computing the following expression under the definition of the variances (6) and (15)

$$\begin{aligned}
&\frac{n_1 n_2}{n}\cdot\left(\frac{1}{\mathrm{Var}(l_1)}+\frac{1}{\mathrm{Var}(l_2)}\right) \\
&= \frac{\frac{n_2}{n}}{p_{T,1}\left(1-p_{T,1}\right)+\frac{n_1}{k_1}p_{C,1}\left(1-p_{C,1}\right)}+\frac{\frac{n_1}{n}}{p_{T,2}\left(1-p_{T,2}\right)+\frac{n_2}{k_2}p_{C,2}\left(1-p_{C,2}\right)} \\
&\to_{n\to\infty} \frac{1-s}{p_{T,1}\left(1-p_{T,1}\right)+t_1 p_{C,1}\left(1-p_{C,1}\right)}+\frac{s}{p_{T,2}\left(1-p_{T,2}\right)+t_2 p_{C,2}\left(1-p_{C,2}\right)}=f
\end{aligned}$$

The convergence follows again by the regularity conditions (16) and (17).

Putting all the pieces together, we get the following representation of χ^2_{net} from (24):

$$\chi^2_{\text{net}} = \left((l_1 - \hat{e}_1)^2 \left(\frac{1}{\hat{v}_1} + \frac{1}{\hat{v}_2} \right) \right) \frac{1}{\hat{w}_n \hat{f}_n}$$

$$= \left(\underbrace{\underbrace{\sqrt{\frac{n}{n_1 n_2}}(l_1 - \hat{e}_1)}_{\to N(0,w)} \cdot \frac{1}{\sqrt{w}}}_{\to N(0,1)} \right)^2 \cdot \underbrace{\frac{\frac{n_1 n_2}{n} \cdot \left(\frac{1}{\text{Var}(l_1)} + \frac{1}{\text{Var}(l_2)} \right)}{\hat{f}_n}}_{\to f/f=1} \cdot \underbrace{\frac{w}{\hat{w}_n}}_{\to 1} \cdot \underbrace{\frac{\frac{1}{\hat{v}_1} + \frac{1}{\hat{v}_2}}{\frac{1}{\text{Var}(l_1)} + \frac{1}{\text{Var}(l_2)}}}_{\to 1}$$

The convergencies follow by (20) to (22) and (25). Thus, χ^2_{net} is asymptotically χ^2-distributed with one degree of freedom which is the distribution of the square of a standard normal distributed random variable.

□

Remark 4.2. In the case of Scenario 2 in Section 3 (the unified target and control groups are representative due to equal target-control rates of the subgroups) the "natural" defintion of $\hat{p}_C = \frac{a_C}{k}$ can be used to define the χ^2_{net}-statistic. It then also follows asymptotically a χ^2-distribution with one degree of freedom, however, the regularity conditions (16) have to be expanded by the assumption that the convergencies are superlinear, i.e. $\lim_{n\to\infty} n\left(\frac{n_1}{n} - s\right) = 0$ and analogously for $\frac{n_1}{k_1}$ and $\frac{n_2}{k_2}$. Also, the latter terms converge to the same number $t_1 = t_2 =: t$. The proof itself becomes more complicated since terms which cancel in the proof of Theorem 4.1 in equations like (23) and (26) only vanish in the limit. In the following we note this version as $\chi^2_{\text{net},1}$.

Remark 4.3. In [5] another slightly different version of χ^2_{net} for Scenario 2, based on $\chi^2_{\text{net},1}$ from Remark 4.2 is presented. It differs since the norming term $\frac{1}{\hat{w}_n \hat{f}_n}$ is omitted and the denominators are defined by

$$\hat{v}_i = n_i \hat{p}_T (1 - \hat{p}_T) + \frac{n_i^2}{k_i} \hat{p}_C (1 - \hat{p}_C) \tag{27}$$

with $\hat{p}_C = \frac{a_C}{k}$. In comparison to (12), the estimation of the variance is based on the whole sample (p_T and p_C) and not only the subgroup specific parts $p_{T,i}$ and $p_{C,i}$. With the above arguments, it can be seen that under the regularity conditions mentioned in Remark 4.2, this statistic also follows a χ^2-distribution with one degree of freedom. However, the null hypothesis has to be expanded in order to include $p_{C,1} = p_{C,2}$, i.e. the hypothesis of equal random noise. In the following we note this version as $\chi^2_{\text{net},2}$.

5. Alternative approaches

In this section we present two alternative methods for approaching the mentioned test scenarios. The first one is the use of contrasts in a variance analytical context shown in [9], the second is a test statistic from net scoring or uplift modelling as presented in [4].

If for an individual observation response is coded as 0 (= no response) and 1 (= response), we can use variance analysis in order to decide our testing problem of equal

uplifts. We regard the group (target or control, subscript: G) as one factor and the two different campaings (or subgroups of one campaign, subscript: SG) as another factor. The empirical means of the four groups are just the estimated probabilities $\hat{p}_{G,SG}$ for a response within the four groups. A classical variance analysis would now compare them in order to search for differences. However, we are interested in the linear hypothesis of the theoretical group means $p_{T,1} - p_{C,1} = p_{T,2} - p_{C,2}$. Such linear hypothosis can be investigated by means of contrasts with the computation of appropriate statistics and p-values, see Section 3.2 of [9]. However, the standard assumptions for variance analysis, namely independence of the observations, normally distributed observations and equal variances in the groups (homoscedasticity), need to be fullfilled.

There are two reasons why we prefer our χ^2_{net}-method. Firstly, "response" and "no response" refer to a binary target variable and by aggregating over the observations we are in the area of count data, for which the χ^2-family of statistics are especially constructed. Variance analysis is primarily aimed at continuous target variables but under certain conditions can be applied to count data.

The second reason emerges from the assumptions of variance analysis. The assumption of independence is standard and is also required for χ^2_{net}. The assumption of normal data can be relaxed when the samples sizes are large, which is usually the case in our marketing applications. However, the third assumption of homoscedasticity is critical in our view. Since the number of responses in each group is binomially distributed, it can be approximated by a normal distribution by the law of large numbers. Different probability parameters for the binomial distribution will automatically result in different variances, therefore spoiling the homoscedasticity condition, compare (6). Since even under the null hypothesis (4) $p_{C,1}$ and $p_{C,2}$ are explicitly allowed to differ from each other, we have heteroscedasticity. Variance analysis is robust to heteroscedasticity when sample sizes are equal, see Section 3.5 of [9]. However, in marketing applications the control group is usually much smaller than the target group (e.g. 10%). Also, Section 3.5 of [9] shows that heteroscedasticity with unequal sample sizes can lead to an increased type I and type II error rate. Thus, we prefer our method which does not suffer from such defects.

The second alternative method is also inspired by a general linear model in combination with a regression. It is described in detail in Section 6.2 of [4]. We will give the formulae here with our notations. Define norming terms by

$$C_{44} := \frac{1}{n_1} + \frac{1}{n_2} + \frac{1}{k_1} + \frac{1}{k_2}$$
$$SSE := n_1\hat{p}_{T,1}(1-\hat{p}_{T,1}) + n_2\hat{p}_{T,2}(1-\hat{p}_{T,2}) + k_1\hat{p}_{C,1}(1-\hat{p}_{C,1}) + k_2\hat{p}_{C,2}(1-\hat{p}_{C,2})$$

and the statistic by

$$t^2_{\text{net}} := \frac{(n+k-4)(\hat{p}_{T,1} - \hat{p}_{C,1} - (\hat{p}_{T,2} - \hat{p}_{C,2}))^2}{C_{44} \cdot SSE} \tag{28}$$

It has the notation t^2_{net} since it is implied (although neither explicitly stated nor proved) in [4] that t_{net} follows asymptotically a t-distribution with $n+k-4$ degrees of freedom. Since $n+k-4$ is quite large in our applications, the t-distribution can be approximated by the standard normal distribution. Thus, t^2_{net} also follows asymptotically a χ^2-distribution with one degree of freedom. In the simulations that follow in the next section, we will see that this statement seems to be true when testing for the right null hypothesis.

6. Simulation study

In this section, we use a simulation study in order to compare the five approaches presented above:

1. χ^2_{net} with p-value $p_{\chi^2_{\text{net}}}$
2. $\chi^2_{\text{net},1}$ with modifications from Remark 4.2 with p-value $p_{\chi^2_{\text{net},1}}$
3. $\chi^2_{\text{net},2}$ with modifications from Remark 4.3 with p-value $p_{\chi^2_{\text{net},2}}$
4. the contrast approach with p-value p_{con}
5. t^2_{net} with p-value $p_{t^2_{\text{net}}}$

For each of the following simulations we consider a fixed set of numbers n_1, n_2, k_1, k_2 and probabilities $p_{T,1}$, $p_{T,2}$, $p_{C,1}$, $p_{C,2}$. The $a_{T,1}$, $a_{T,2}$, $a_{C,1}$ and $a_{C,2}$ are binomial $B(n_1, p_{T,1})$-, $B(n_2, p_{T,2})$-, $B(k_1, p_{C,1})$- and $B(k_2, p_{C,2})$-distributed random variables. With these our basic stochastic model as shown in Section 4 is completely described. Being binomially distributed, $a_{G,SG}$ are easily simulated with any standard statistical software package. Since all five statistics and p-values above are functions of the n_i, k_i and $a_{G,SG}$, they and their corresponding p-values can be computed. For a fixed set of parameters, we repeat this b times, usually choosing $b = 100$. We then sort the corresponding data set by the p-values $p_{\chi^2_{\text{net}}}$ and denote the resulting numbers by $p^{i:b}_{\chi^2_{\text{net}}}$, $p^{i:b}_{\chi^2_{\text{net},1}}$, $p^{i:b}_{\chi^2_{\text{net},2}}$, $p^{i:b}_{\text{con}}$, $p^{i:b}_{t^2_{\text{net}}}$, $i = 1, \ldots, b$. Remark that only the $p^{i:b}_{\chi^2_{\text{net}}}$ are necessarily in ascending order for $i = 1, \ldots, b$. We then plot each of the five p-value series against the set of $\frac{i}{b}$, $i = 1, \ldots, b$, i.e. the points $\left(\frac{i}{b}, p^{i:b}_x\right)$ for $x = \chi^2_{\text{net}}, \chi^2_{\text{net},1}, \chi^2_{\text{net},2}, \text{con}, t^2_{\text{net}}$. If the null hypothesis behind each test is fulfilled, the p-values follow a uniform distribution on $(0, 1)$ and the plotted points scatter around the diagonal in this probability plot. However, if the null hypothesis is not fulfilled, we will find deviations from the diagonal.

Table 1 shows the parameter values behind the results in Figures 1 to 7.

	aimed at error	n_1	n_2	k_1	k_2	$p_{T,1}$	$p_{T,2}$	$p_{C,1}$	$p_{C,2}$
Figure 1	type I	50,000	50,000	5,000	5,000	10%	10%	9%	9%
Figure 2	type I	50,000	50,000	5,000	5,000	5%	51%	4%	50%
Figure 3	type I	100,000	20,000	10,000	2,000	5%	51%	4%	50%
Figure 4	type I	100,000	20,000	10,000	2,000	51%	5%	50%	4%
Figure 5	type II	50,000	50,000	5,000	5,000	11%	10%	9%	9%
Figure 6	type II	50,000	50,000	5,000	5,000	6%	51%	4%	50%
Figure 7	type II	50,000	50,000	5,000	10,000	5%	52%	4%	50%

Table 1. Parameters used for the seven simulations.

The control groups were chosen to be roughly 10% of the target group which is a quite common target-control rate. Also, the absolute group sizes are not unusual in practice.

In general, all figures show one fact: there are very small differences between p_{con} and $p_{t^2_{\text{net}}}$. Thus, in essence, although not clear from the description in Section 5 they seem to be the same method. This means that all the criticism stated for p_{con} carries over to $p_{t^2_{\text{net}}}$. The points of criticism are supported by the simulations.

In Figures 1 to 4 we investigate if the tests control the type I error, since the uplift in both subgroups is 1%. We vary the rate between the target groups and the levels of the random noise in the subgroups. The results are as follows:

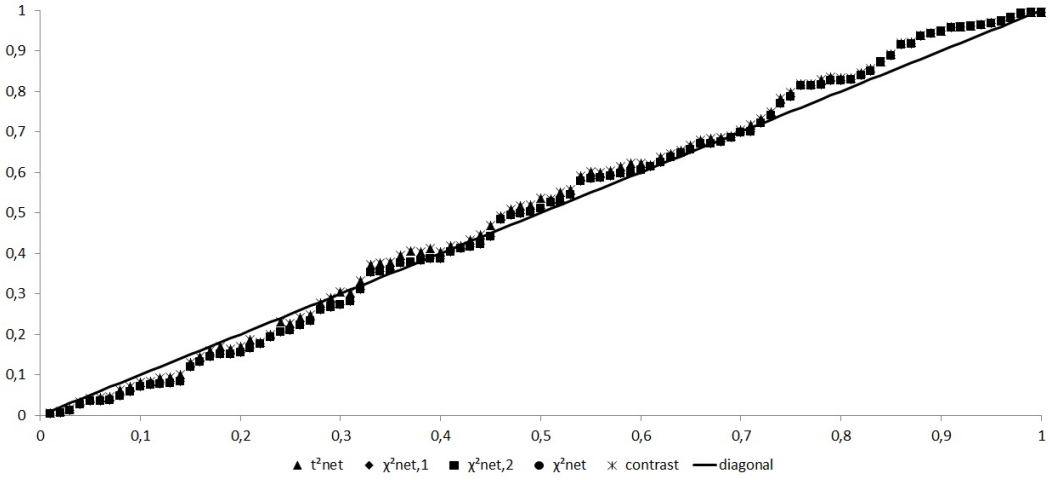

Figure 1. Check for type I error with equal group sizes and equal random noise.

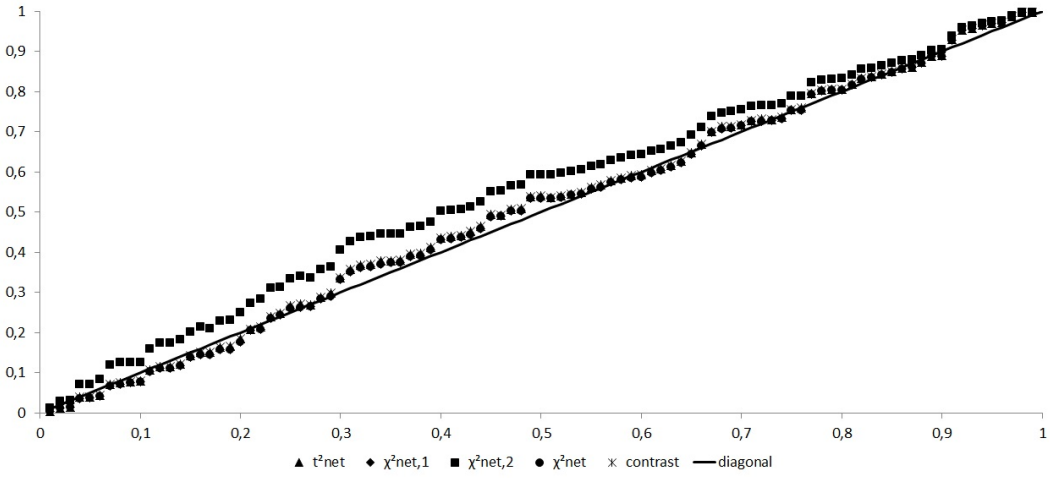

Figure 2. Check for type I error with equal group sizes and different random noise.

- When both target groups have roughly the same size and the random noise is equal, all methods deliver nearly the same results and maintain the type I error (Figure 1).
- When target group sizes are the same, however random noises are on a different scale, $\chi^2_{\text{net},2}$ gives slightly different results, however, the type I error is still kept by all methods (Figure 2).
- When target group sizes differ and the smaller group has the larger random noise, only χ^2_{net} and $\chi^2_{\text{net},1}$ are able to keep the type I error. t^2_{net} has the smallest p-values and is no reliable in this case (Figure 3).
- When target groups differ heavily and the larger group has the larger random noise, only χ^2_{net} and $\chi^2_{\text{net},1}$ are able to control the type I error. The other methods have p-values too high and are not reliable tests here (Figure 4).

In Figures 5 to 7, we investigate the type II error with an uplift of 2% in one subgroup and of 1% in the other subgroup. Further, we only look at the case of equal target group sizes, since in the other case only χ^2_{net} and $\chi^2_{\text{net},1}$ can reasonably be used by the above results. Here we see:

- When random noise is roughly the same size, all methods detect deviations from the null hypothesis with χ^2_{net} showing the smallest p-values (Figure 5).

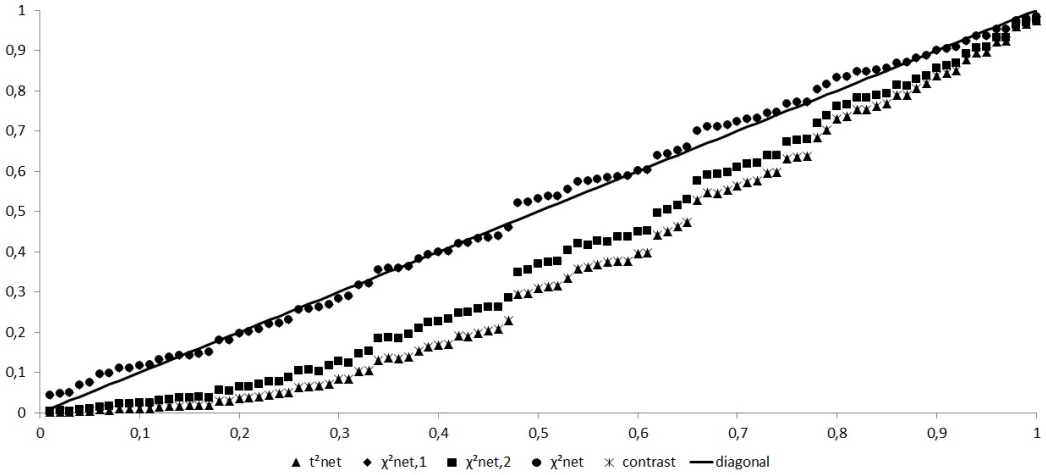

Figure 3. Check for type I error with different group sizes and different random noise (small group with large noise).

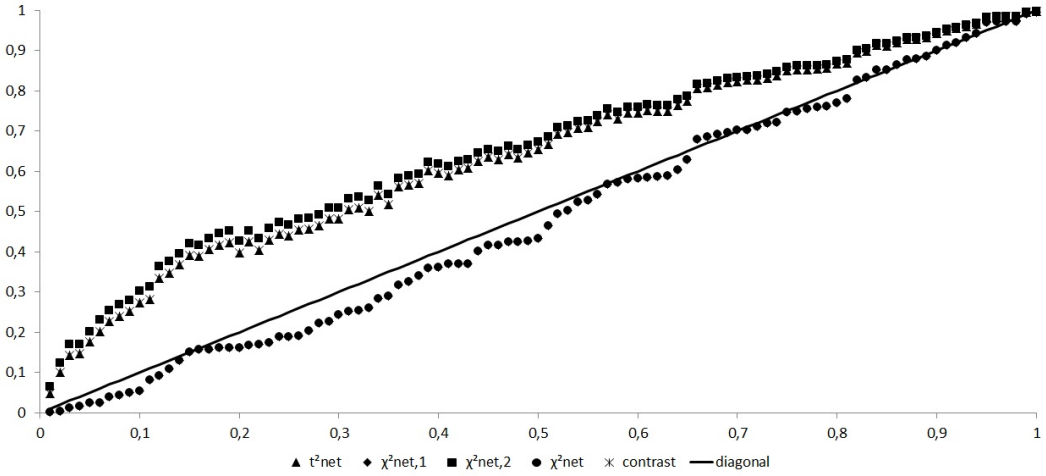

Figure 4. Check for type I error with different group sizes and different random noise (large group with large noise).

- When random noises are on a different scale, $\chi^2_{\text{net},2}$ has a notably larger type II error, the others are roughly equal (Figure 6).
- When the target-control rate differs between the subgroups, $\chi^2_{\text{net},1}$ and $\chi^2_{\text{net},2}$ are useless and not shown since they were not constructed for this case. Of the remaining methods, χ^2_{net} is clearly the best one with the smallest type II error (Figure 7).

The results of the simulations are quite clear: χ^2_{net} should be the method of choice when testing for equal uplifts. It seems that t^2_{net} and the contrast method not only test for the null hypothesis of equal uplifts but additionally assume equal random noise, as does $\chi^2_{\text{net},2}$, see Remark 4.3. A mathematical investigation and comparison of the type II error rates of all methods, but especially χ^2_{net} and t^2_{net}, is still an open issue.

7. Application to real data

We next want to apply the statistics to real world data, in particular covering both scenarios from Section 3, i.e. one campaign is split up or two different campaigns are joint together.

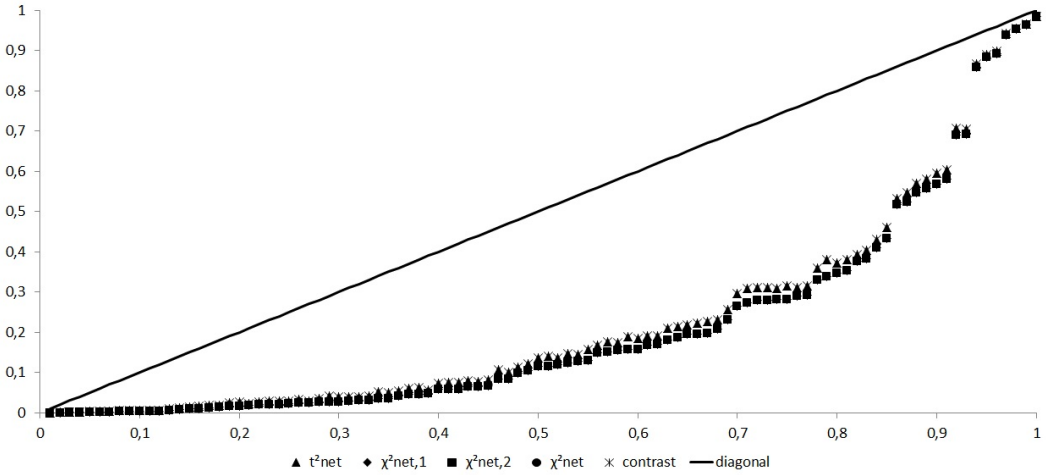

Figure 5. Check for type II error with equal group sizes and equal random noise.

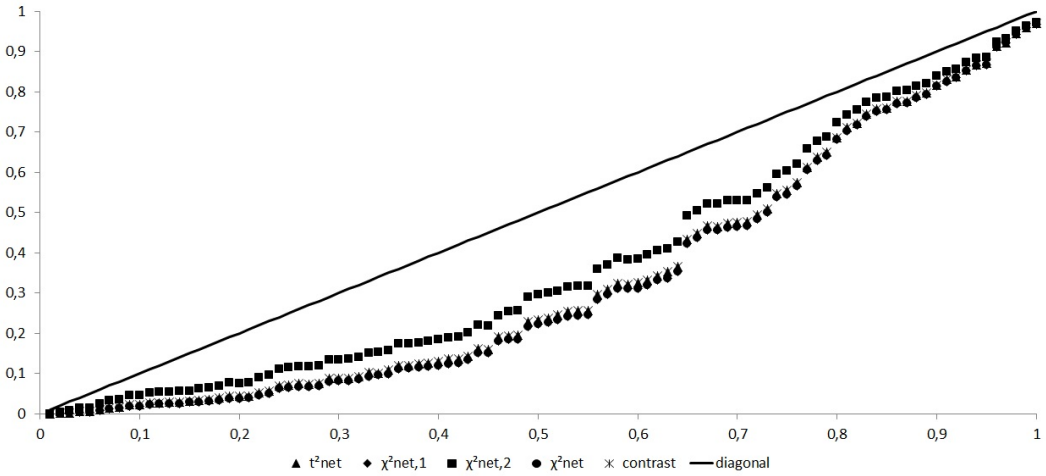

Figure 6. Check for type II error with equal group sizes and different random noise.

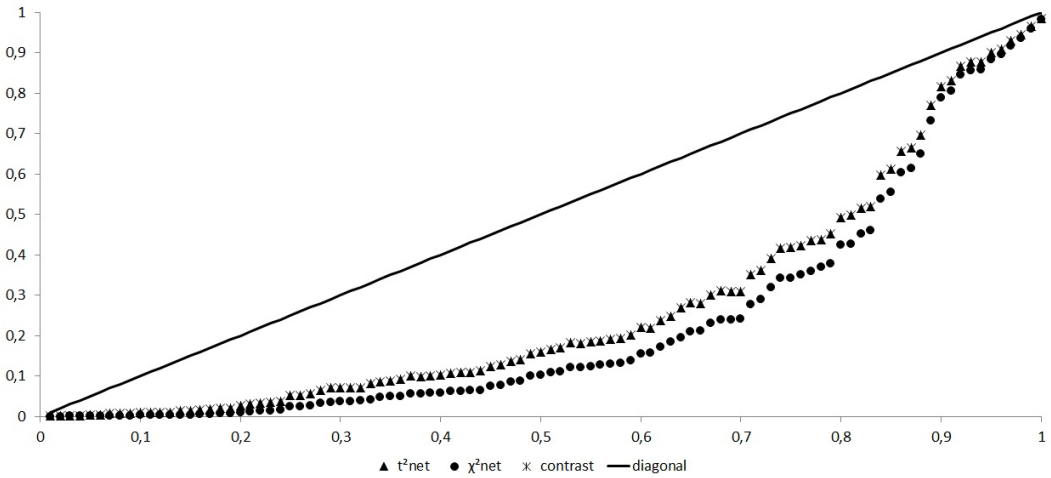

Figure 7. Check for type II error with equal group sizes and different random noise under different target-control rate.

A campaign was done by a financial institute with the aim of generating appointments with regard to provision for one's old age. For that purpose, 167 thousand customers were selected to receive a letter and a phone call inviting them to make an appointment with their bank advisor in order to discuss their hedging situation in old age. A response was counted if an appointment was made. A control group of 13 thousand customers was created. Accordingly, appointments were counted as responses.

The first thing of interest for the bank was if gender is a decisive factor in response to the campaign. The bank had the hypothesis that women are more affine to provision for old age since they tend to think more about saving for the future than men who are supposed to spend money more quickly.

Table 2 shows the results of the campaign, seperated by gender.

campaign 1	target			control			uplift
	persons	responses	rate	persons	responses	rate	
women	81,770	5,656	6.92%	6,391	373	5.84%	1.08%
men	85,257	6,231	7.31%	6,699	443	6.61%	0.70%

Table 2. Comparison of campaign results by gender.

The results seem to confirm the stronger affinity of women to provision, since women had an uplift of 1.08% and men one of 0.70%. A classical χ^2-test for homogenity shows for men and women that the response rates in target and control group are significantly different from each other in both groups, i.e. the uplifts are significantly different from 0. The question is, however, if the uplifts are statistically different from each other. In order to answer this question, we compute $\chi^2_{\text{net}} = 0.7643$ and the corresponding p-value $p_{\chi^2_{\text{net}}} = 0.3820$. Thus, in this campaign, the uplifts of women and men are not significantly different and the hypothesis above cannot be confirmed. Since we are in Scenario 2 ($\frac{n_1}{k_1} = \frac{81{,}770}{6{,}391} \approx 12.79 \approx 12.73 \approx \frac{85{,}257}{6{,}699} = \frac{n_2}{k_2}$), the two alternatives from Remarks 4.2 and 4.3 can also be computed. They result in $\chi^2_{\text{net},1} = 0.7648$ and $\chi^2_{\text{net},2} = 0.7622$ with corresponding p-values of $p_{\chi^2_{\text{net},1}} = 0.3818$ and $p_{\chi^2_{\text{net},2}} = 0.3827$. Thus, the alternatives lead to almost exactly the same results. The approach by means of contrasts generates the p-value $p_{\text{con}} = 0.4075$, i.e. no significance. The statistic $t^2_{\text{net}} = 0.6861$ with the same p-value of $p_{t^2_{\text{net}}} = 0.4075$ leads to the same conclusion.

Parallel to Campaign 1, another campaign was carried out which had the same objective (to get customers to make an appointment with their bank advisor to talk about old age provision) and the same means (firstly a letter was sent and afterwards, some of the customers were called), however, the letter had a different layout and text. Also, the customers were different since the second campaign aimed at customers with higher income. Was the effort more successful for wealthier clients? Table 3 shows the results.

A classical χ^2-test once again shows that the campaigns themselves were successful since both uplifts are significantly different from 0. When comparing the uplifts to each other, the uplift for the wealthier clients is 1.61% which is above the uplift of 0.88% for the middle-class clients. In order to check for statistical significance, we compute $\chi^2_{\text{net}} = 3.8661$ and $p_{\chi^2_{\text{net}}} = 0.0493$ which is significant at the usual 5%-level. Thus, there is evidence that a campaign with the above targets is more effective for customers with higher income. Since $\frac{n_1}{k_1} = \frac{167{,}027}{13{,}090} \approx 12.76 \gg 5.55 \approx \frac{44{,}356}{7{,}987} = \frac{n_2}{k_2}$, we are in Scenario 1 and $\chi^2_{\text{net},1}$ and $\chi^2_{\text{net},2}$ cannot be applied.

However, the alternative approach by means of contrasts can be applied. The p-value here is $p_{\text{con}} = 0.0626$ and, thus, not significant, although close to significant. t^2_{net}-statistic

	target			control			uplift
	persons	responses	rate	persons	responses	rate	
campaign 1	167,027	11,887	7.12%	13,090	816	6.23%	0.88%
campaign 2	44,356	3,447	7.77%	7,987	492	6.16%	1.61%

Table 3. Comparison of results for parallel campaigns aimed at different customer segments.

gives the value of 3.4672 with a p-value of $p_{t^2_{\text{net}}} = 0.0626$, thus showing no significance in contrast to χ^2_{net}. This is another indication that the statistical power of the t^2_{net} or the contrast method is lower than that of χ^2_{net}. This example too shows that both methods (t^2_{net} and contrast) coincide at least quantitatively.

8. Discussion and outlook

In this article, we have presented a new statistic, based on the classical χ^2 which is appropriate when statistically comparing the uplift of two campaigns. We have proved that its asymptotic distribution is a χ^2-distribution with one degree of freedom and shown its practical applicability by using it on real data to decide a real life problem. We have also shown by means of simulation that it seems to be superior to the already existing alternative approaches by fixing the type I error and showing smaller type II errors.

However, an open issue remains in the comparison of the method presented here with the alternative ones by mathematical means.

The statistic χ^2_{net} has been presented here for the two sample (campaign) case. A generalization of the formula of χ^2_{net} to the $j > 2$ sample case seems straightforward by

$$\chi^2_{\text{net}} := \frac{1}{\hat{w}_n \hat{f}_n} \sum_{i=1}^{j} \frac{(l_i - \hat{e}_i)^2}{\hat{v}_i}$$

However, the definition of suitable norming terms $\hat{w}_n$ and $\hat{f}_n$ is yet unclear. The conjecture is, of course, that this generalized version is asymptotically distributed as the χ^2 distribution with $j-1$ degrees of freedom. A mathematical assessment of that assertion still needs to be done.

For the classical χ^2-statistic, rules of thumb are known when the approximative distribution is valid (e.g. the expected frequency in each cell must be larger than 5, i.e. $np_i > 5$, see Section 4.2 of [6]). Such rules are still missing in the χ^2_{net}-case.

Besides its applications to testing problems in marketing performance measurement, this statistic can also be used a scoring context, the so called net scoring. Especially its application to the construction of decision trees was described in [5].

However, marketing is not the only area in which our statistic can be applied. It is useful in all testing scenarios where the effect of a treatment is investigated under the condition that the desired result could also appear by itself. The medical application seems to suggest itself since typically the impact of different drugs is compared well-knowing that a certain percentage of patients will recover even without treatment. Examples of the net effect in medicine are shown in [10] and [11]. We hope this article stipulates research in this area and the discovery of many more areas of application.

Acknowledgements

The authors would like to thank Hans Fischer, Michael Falk and Johannes Hain for their support on this work as well as Altran GmbH & Co. KG.

References

[1] Radcliffe N, Simpson R. Identifying who can be saved and who will be driven away by retention activity. Journal of Telecommunications Management. 2008;1(2):168–176.
[2] Lo V. The true lift model - a novel data mining approach to response modeling in database marketing. SIGKDD Explorations. 2002;4(2):78–86.
[3] Falk M, Fischer H, Hain J, Marohn F, Michel R. Statistik in theorie und praxis - mit anwendungen in r. Munich: Springer; 2014. to appear.
[4] Radcliffe N, Surry P. Real-world uplift modeling with significance-based uplift trees. 2011; Technical Report, Stochastic Solutions.
[5] Michel R, Schnakenburg I, von Martens T. Methods of variable pre-selection for netscore modeling. Journal of Research in Interactive Marketing. 2013;7(4):257–268.
[6] Falk M, Marohn F, Tewes B. Foundations of statistical analyses - examples with sas. Basel: Birkhäuser; 2003.
[7] Schinazi R. Probability with statistical applications. 2nd ed. Birkhäuser; 2012.
[8] Cramér H. Mathematical methods of statistics. Princeton: University Press; 1945.
[9] SAS. Statistics 2: Anova and regression course notes. Cary: SAS Institute Inc.; 2012.
[10] Jaskowski M, Jaroszewicz S. Uplift modeling for clinical trial data. ICML 2012 Workshop on Clinical Data Analysis. 2012;.
[11] Nassif H, Kuusisto F, Burnside E, Page D, Shavlik J, Santos Costa V. Score as you lift (sayl): A statistical relational learning approach to uplift modeling. Proceedings of the European Conference on Machine Learning (ECML). 2013;.
[12] Hansotia B, Rukstales B. Incremental value modeling. Journal of Interactive Marketing. 2002;16(3):35–46.
[13] Rzepakowski P, Jaroszewicz S. Decision trees for uplift modeling with single and multiple treatments. Knowledge and Information Systems. 2011;28(2):303–327.